

Optimal Mix of Incentive Strategies for Product Marketing on Social Networks

Pankaj Dayama¹, Aditya Karnik¹ and Y. Narahari²

¹ Global General Motors R&D, India Science Lab, GM Technical Center (India), Bangalore 560066, India

{pankaj.dayama, aditya.karnik}@gm.com

² Department of Computer Science and Automation, Indian Institute of Science, Bangalore 560012, India
hari@csa.iisc.ernet.in

Abstract. We consider the problem of devising incentive strategies for viral marketing of a product. In particular, we assume that the seller can influence penetration of the product by offering two incentive programs: a) direct incentives to potential buyers (*influence*) and b) referral rewards for customers who influence potential buyers to make the purchase (*exploit connections*). The problem is to determine the optimal timing of these programs over a finite time horizon. In contrast to algorithmic perspective popular in the literature, we take a mean-field approach and formulate the problem as a continuous-time deterministic optimal control problem. We show that the optimal strategy for the seller has a simple structure and can take both forms, namely, *influence-and-exploit* and *exploit-and-influence*. We also show that in some cases it may optimal for the seller to deploy incentive programs mostly for low degree nodes. We support our theoretical results through numerical studies and provide practical insights by analyzing various scenarios.

Keywords: Incentive Strategies, Social Networks, Viral Marketing

1 Introduction

A key research topic in multi-agent systems is to understand the effect of microdynamics/interactions between agents on macroscopic properties. Often the agents are a part of a social structure such as a social network. A common example is that of a social network that consists of potential buyers of a particular new product offering in the market. These buyers interact with each other and influence each others' purchase decisions through word-of-mouth and/or behavior. This so-called *social influence* exerted by agents on their neighboring agents in the network have a significant role to play in generating a network effect on the sales of the product. The idea of viral marketing is to essentially exploit the (macroscopic) network effects that result due to the microdynamics between the agents in the network.

Viral marketing is receiving much attention by practicing marketers and academics alike. While not a new idea, it has come to the forefront because of multiple effects - products have become more complex, making buyers to increasingly rely on opinions of their peers; consumers have evolved to distrust advertising; and Web2.0 has revolutionized the way people can connect, communicate and share. With power shifting to consumers, it has become important for sellers to devise effective viral marketing strategies [8]. This work is motivated by this urgent need.

For social influence to work, there must be *seeds*, i.e., product advocates to start with. The sellers, therefore, employ two basic strategies. The first is to create advocates, by providing incentives to potential buyers to make an actual purchase. These incentives are typically in the form of discounts, free goodies, etc. The second is to reward product advocates who ‘put in a good word’ and influence potential buyers to make the purchase. Thus, the latter program helps to *exploit* the impact of social influence while making a purchasing decision whereas the former program helps to directly *influence* the buying behavior by offering discounts.

Since incentives come at a cost, a seller must balance the revenue she generates through these strategies and the expenditure she incurs in doing so. This poses some non-trivial challenges. The first is determining incentives themselves, since response of an individual is contingent on them (too low a referral reward may not elicit recommendation from an individual since personal reputation is usually at stake). Secondly, the two programs are not necessarily causally connected. The reputation of a firm or a brand might create product advocates without incentives, thereby, requiring a seller to launch a referral program directly. This necessitates careful ‘timing’ of these programs.

The objective of this paper is to shed some light on this practically important and theoretically interesting problem. In particular, we seek to determine an optimal timing of these programs over a finite time horizon.

1.1 Related Work

In recent years, problems such as these have attracted much attention. Several papers investigate ‘influence maximization’ (see, for example, Domingos and Richardson [7], Kempe, Kleinberg, and Tardos [11], Bharathi, Kempe, and Salek [4], Chen, Yaun, and Zhang [6]), where the problem is to determine the set of initial adopters who, through an influence process, can maximize the future adoptions of the product. Auriol and Benaim [2] discuss a dynamic model of how standards and norms emerge in decentralized economies. Hartline, Mirrokni, and Sundararajan [10], Arthur, Motwani, Sharma, and Xu [1] consider the problem of ‘revenue maximization’ for viral marketing and are close in spirit to the problem we consider in this paper.

In Hartline et al. [10] a model is proposed in which the purchase decision of a buyer is influenced by individuals who own the product and the price at which the product is offered. An optimal pricing policy is derived using dynamic programming in a symmetric setting (i.e., identical buyers). In a general setting,

finding an optimal strategy is shown to be NP-hard and approximation algorithms are considered. The authors suggest *influence-and-exploit* strategy where selected buyers are given the product for free, and the seller extracts revenue by making a random sequence of offers and a greedy pricing strategy for the remaining buyers to compensate for the initial loss.

Arthur et al. [1] also considers a model in which a buyer's decision is influenced by friends who own the product and price at which the product is offered. Sales are assumed to cascade through the social network. The seller offers cash-back to recommenders and also sets price for each buyer. The authors show that determining an optimal strategy to maximize expected revenue is NP-hard and propose a non-adaptive *influence-and-exploit* policy, which offers product to the interior nodes of the max-leaf spanning tree of the network for free and later exploits their influence by extracting more revenue from the leaf nodes of the tree. They show that the expected revenue generated from the non-adaptive strategy is within a constant factor of the optimal revenue from an adaptive strategy.

1.2 Our Contributions

We consider a seller interested in selling a product to a population of N agents. The product is assumed to be durable and free from network externalities. From the seller's perspective, each agent assumes one of the following types at any point in time: *potential buyer* (one who is yet to make a purchase), *customer* (one who has purchased the seller's product) and *competitor's customer* (one who has purchased a competing product³). A potential buyer makes a purchase decision of her own volition (essentially under *external influence*) or under social influence. This decision-making is modeled probabilistically, by specifying for both products (the seller's and the competitor's), probabilities of purchase under external influence and social influence. The seller can influence the former through direct incentives (which affect the price) and the latter through referral rewards. The problem she faces is to roll out these programs so as to maximize the profit, which is equal to the revenue obtained by customer acquisition minus the expenditure on direct incentives and referral rewards, over a given time horizon T .

In practice sellers have limited knowledge about the social network underlying the population, typically, in the form of a *class*-level statistical description of it. A class comprises agents who are considered essentially identical on a variety of factors (chosen by the seller), such as demographic, economic level, number of social contacts and so on. In this paper we consider this set-up. However to keep it simple, we assume heterogeneity only in terms of network connections (in particular, probabilities of purchase under either external or social influence are assumed to be the same for all agents); hence classes are based only on the number of social contacts (degrees). The seller thus knows only the degree distribution and degree-degree correlation of the social network.

³ All competitors are aggregated into one single virtual competitor.

This class-level statistical description of the agent population allows us to approximate the stochastic evolution of the purchase dynamics by a deterministic process described by ordinary differential equations (ODEs). This is formally established as a mean-field limit, taking the number of agents $N \rightarrow \infty$ [3]. With an ODE limit, we pose the problem as a continuous time optimal control problem and employ the well known Pontryagin's Maximum principle [12] to characterize an optimal control. An optimal control specifies for each class the times at which direct incentives and referral rewards programs are to be executed. The following are our main results.

1. We show that an optimal control has a simple structure: the seller needs to run each of the programs at most twice for a certain duration. Moreover, it is non-adaptive (or open-loop). This simplifies the implementation and practically can help a seller pre-allocate the budget for her campaign.
2. While *influence-and-exploit* strategy turns out to be optimal when social influence is strong in the population, *exploit-and-influence* strategy can be optimal when the seller has a good reputation.
3. In some settings, the seller may be better off incentivizing low degree nodes as against the popular approach of targeting the influentials (high degree nodes). This, we believe, provides some support to the findings reported in [14] in reference to the *influentials hypothesis*.

The approach we have taken to address the problem is entirely different from the ones in the literature. While a large size of the population presents a challenge to the earlier approaches, it, in fact, aids us in migrating to a simpler deterministic description of the dynamics. The assumption that agents of a class are indistinguishable also fits in naturally with the popular marketing approach of *customer segmentation* and allows a seller to customize incentives and referral rewards as per these segments.

In contrast to earlier papers, we have also modeled competition. This is not only close to reality but interestingly it allows to address some problems in completely different contexts. For example, in limiting the spread of misinformation about an entity or an Internet virus, the objective is to maximize nodes with correct information or security patches by immunizing them (akin to direct incentives) and/or incentivizing them to spread the information they have to their neighbors (akin to referrals). Our results are, thus, applicable to these problems as well (see [5] for discussion of the influence limitation problem).

2 Problem Formulation

Consider a population of N agents, indexed by $i = 1, 2, \dots, N$. The underlying social network is specified by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Each agent is identified with a node in \mathcal{V} and $(i, j) \in \mathcal{E}$ means that i and j are social contacts and they influence each other in decision-making.

s_i denotes the state of agent i . s_i can take three values: 0 (indicates potential buyer), 1 (indicates customer) and -1 (indicates competitor's customer). Let $\mathbf{s} := (s_1, s_2, \dots, s_N)$.

Each agent makes the purchase decision at a random time point, independent of all others. It suffices to assume that time is discrete (denoted by $n = 1, 2, \dots$) and at each time step, an agent is chosen uniformly randomly from the population for a potential state change. Since there are no repeat purchases, 1 and -1 are absorbing states. Therefore, the state change occurs only if the chosen agent is a potential buyer. Suppose agent i is chosen at a time n . Then one of the following happens if i is a potential buyer.

1. i buys the seller's product on her own with probability α (For example, $\alpha = 0.08$ means that there is 8% chance that a potential buyer will buy the seller's product on her own).
2. i buys the competitor's product on her own with probability δ .
3. i selects one of her social contacts at random. If the selected contact is a customer, i buys the seller's product under social influence with probability β (For example, $\beta = 0.1$ means that there is 10% chance that a potential buyer will buy the seller's product if she interacts with someone who has already bought the product).
4. i selects one of her social contacts at random. If the selected contact is a competitor's customer, i buys the competitor's product under social influence with probability γ .

Clearly, the state process $\{\mathbf{s}(n), n \geq 1\}$ is a Markov chain.

Now from the seller's perspective, agents having the same degree are indistinguishable and the network \mathcal{G} is known only statistically, i.e., \mathcal{G} is drawn from an ensemble of random undirected graphs of size N , a given degree distribution $P(k)$ ($1 \leq k \leq K$) and degree-degree correlation function $P(k'|k)$, which denotes the probability that a given link from a node of degree k is to a node of degree k' . Note that a number of well-known graphs such as homogeneous random graphs, exponential random graphs (e.g., $G(n, p)$ and Watts-Strogatz network), scale-free networks can be represented in this framework. We assume that K remains uniformly bounded as $N \rightarrow \infty$.

Denote by i_k , r_k and θ_k the fraction of degree- k agents who are potential buyers, customers and competitor's customers respectively (note that normalization is with respect to the number of class- k agents; hence $i_k + r_k + \theta_k = 1$). Let $x_k := (i_k, r_k, \theta_k)$ and $\mathbf{x} := (x_1, \dots, x_K)$. From the above assumption, it follows that $\{\mathbf{x}(n), n \geq 1\}$ is a Markov chain (as seen by the seller).

The drift of \mathbf{x} can be computed considering the four cases described above. Table 1 shows the corresponding probabilities and the change in x_k for degree

class-k. Consider as an example Case 3. The probability of a randomly selected agent being a potential customer of degree k is $P(k)i_k$. This agent randomly chooses one of her k social contacts. The probability that this chosen one is an existing customer is $\frac{1}{k} \sum_{j=1}^k \sum_{k' \in K} P(k'|k)r_{k'} = \sum_{k' \in K} P(k'|k)r_{k'}$. The selected agent buys the seller's product under the social influence from her contact with probability β . Thus the probability of Case 3 is $\beta P(k)i_k \sum_{k' \in K} P(k'|k)r_{k'}$. One agent changes her state from 0 (potential customer) to 1 (customer). Hence the effect on x_k is $\frac{1}{NP_k}(-1, 1, 0)$.

We now make the dependence on the population size N explicit and denote by $F^N(\mathbf{x}) := [F_k^N(\mathbf{x})]_{k=1}^K$ the drift of \mathbf{x} . $F_k^N(\mathbf{x})$ is as follows.

$$\frac{1}{N} \begin{pmatrix} -\beta i_k \sum_{k' \in K} P(k'|k)r_{k'} - \gamma i_k \sum_{k' \in K} P(k'|k)\theta_{k'} - (\alpha + \delta)i_k \\ \beta i_k \sum_{k' \in K} P(k'|k)r_{k'} + \alpha i_k \\ \gamma i_k \sum_{k' \in K} P(k'|k)\theta_{k'} + \delta i_k \end{pmatrix}$$

Observe that (i) the number of transitions per agent per time slot is of the order of $\frac{1}{N}$ (ii) the second moment of number of agent transitions per time slot is bounded and (iii) $F^N(\mathbf{x})$ is a smooth function of $\frac{1}{N}$ and \mathbf{x} . Let $F(\mathbf{x}) = \lim_{N \rightarrow \infty} \frac{F^N(\mathbf{x})}{1/N}$. It then follows from Theorem 1 of [3] that the time evolution of $\mathbf{x}(n)$ can be approximated by the following system of ODEs (with the same initial conditions).

$$\dot{\mathbf{x}} = F(\mathbf{x}) \quad (1)$$

More explicitly, for $1 \leq k \leq K$

$$\begin{aligned} \dot{i}_k &= -\beta i_k R_k - \gamma i_k \Theta_k - (\alpha + \delta)i_k \\ \dot{r}_k &= \beta i_k R_k + \alpha i_k \\ \dot{\theta}_k &= \gamma i_k \Theta_k + \delta i_k \end{aligned}$$

where, $R_k := \sum_{k' \in K} P(k'|k)r_{k'}$ and $\Theta_k := \sum_{k' \in K} P(k'|k)\theta_{k'}$.

The seller offers direct incentives and referral rewards to increase α and β respectively. We model this as follows. A referral reward of c results in an increase of ϵ_1 in β and a direct incentive of c' causes α to increase by ϵ_2 . Thus, for the duration of the referral reward program, social influence rate of $(\beta + \epsilon_1)$ is operational and the seller incurs a cost of c for every successful referral. Similarly,

Case	Probability	Effect on x_k
1	$\alpha P(k)i_k$	$\frac{1}{NP_k}(-1, 1, 0)$
2	$\delta P(k)i_k$	$\frac{1}{NP_k}(-1, 0, 1)$
3	$\beta P(k)i_k \sum_{k' \in K} P(k' k)r_{k'}$	$\frac{1}{NP_k}(-1, 1, 0)$
4	$\gamma P(k)i_k \sum_{k' \in K} P(k' k)\theta_{k'}$	$\frac{1}{NP_k}(-1, 0, 1)$

Table 1. Probability and effect on x_k for different cases

if the direct incentive program is executed for some duration, the take-rate for seller's product increases to $(\alpha + \epsilon_2)$ for that duration, incurring her a cost of c' for every sale. We normalize c and c' with respect to the product price. Thus the price is fixed to 1. The seller's problem of maximizing her profit (revenue minus cost) over a fixed time horizon T by optimally timing the two program can now be stated formally as follows.

Let $u_k(t)$ (resp. $v_k(t)$) denote the control variable indicating whether or not the referral reward program (resp. direct incentive program) is offered to class-k at time t . The cost incurred in running the referral reward program is

$$\int_0^T \sum_{k=1}^K P(k) u_k(t) c(\beta + \epsilon_1) i_k(t) R_k(t) dt \quad (2)$$

Recall that the *conversion rate* of potential buyers under the program is $c(\beta + \epsilon_1) i_k(t) R_k(t)$. The cost incurred in the direct incentives program is

$$\int_0^T \sum_{k=1}^K P(k) v_k(t) c'(\alpha + \epsilon_2) i_k(t) dt \quad (3)$$

Since the product price is unity, the revenue obtained is proportional to the number of customers at the end of horizon, $\sum_{k=1}^K P(k) r_k(T)$. Denoting the total cost (2)+(3) by $C(T)$ the problem is

$$\text{Maximize } \sum_{k=1}^K P(k) r_k(T) - C(T)$$

subject to

$$\begin{aligned} \dot{i}_k &= -(\beta + u_k \epsilon_1) i_k R_k - (\alpha + v_k \epsilon_2) i_k - \gamma i_k \Theta_k - \delta i_k \\ \dot{r}_k &= (\beta + u_k \epsilon_1) i_k R_k + (\alpha + v_k \epsilon_2) i_k \\ \dot{\theta}_k &= \gamma i_k \Theta_k + \delta i_k \end{aligned}$$

for $1 \leq k \leq K$ and a given initial condition $\mathbf{x}(0)$.

Three remarks are in order. The assumption of heterogeneity only in the number of social contacts is mainly to keep the formulation simple and highlight the impact of network structure. Extending this formulation to a general setting is straightforward and will be taken up in a longer version of the paper. Our random interaction model essentially means that the social influence on a potential buyer is the average influence from her neighbors. This, we believe, is reasonable since we have also assumed presence of external influence (through α) on agents⁴. In the above formulation we consider fixed rewards and incentives pay-outs (c and c'). This simplifies implementation in practice. Calibration of c and c' can be carried out through numerical studies.

⁴ For the lack of clear empirical evidence, one may also consider *total* influence from the neighbors. Mathematically, it is a simple modification to our formulation.

3 Structure of Optimal Control

In this section we mathematically prove the structural properties of an optimal control. To keep the proof simple, we will assume that the network \mathcal{G} is drawn randomly from a set of regular networks of size N and degree k . This is without loss of generality.

Let $i(t)$, $r(t)$ and $\theta(t)$ denote the fraction of population in states $\{0, 1, -1\}$ at time t respectively. Let $u(t) \in \{0, 1\}$ denote whether or not the referral reward program is offered at time t and let $v(t) \in \{0, 1\}$ denote whether or not the direct incentive program is offered at time t . The purchase dynamics under the influence of these programs are given as follows:

$$\dot{i} = -(\beta + u\epsilon_1)ir - (\alpha + v\epsilon_2)i - \gamma i\theta - \delta i \quad (4)$$

$$\dot{r} = (\beta + u\epsilon_1)ir + (\alpha + v\epsilon_2)i \quad (5)$$

$$\dot{\theta} = \gamma i\theta + \delta i \quad (6)$$

From (4), (5), and (6), observe that $\dot{i} + \dot{r} + \dot{\theta} = 0$. Therefore, it suffices to consider any two equations. Let $\Omega := \{(i, r) | i + r \leq 1, i \geq 0, r \geq 0\}$. Let $\mathbf{x}(t) := (i(t), r(t)) \in \Omega$ denote the state variable.

The optimal control problem in this simpler setting is as follows.

$$\begin{aligned} \text{Maximize } & r(T) - \int_0^T cu(t)(\beta + \epsilon_1)i(t)r(t)dt \\ & - \int_0^T c'v(t)(\alpha + \epsilon_2)i(t)dt \end{aligned} \quad (7)$$

subject to (4), (5) and the following constraints on state and control variables: for all $0 \leq t \leq T$, $\mathbf{x}(t) \in \Omega$, $u(t) \in \{0, 1\}$ and $v(t) \in \{0, 1\}$.

Our main result is given in Proposition 1. It shows that an optimal strategy for the seller is to deploy the two incentive programs for at most two distinct time periods.

Proposition 1. 1. There exist τ_1, τ_2 ($0 \leq \tau_1 \leq \tau_2 \leq T$) such that $u^*(t) = 0$ for $\tau_1 < t \leq \tau_2$ and $u^*(t) = 1$ else.
2. There exist τ_3, τ_4 ($0 \leq \tau_3 \leq \tau_4 \leq T$) such that $v^*(t) = 0$ for $\tau_3 < t \leq \tau_4$ and $v^*(t) = 1$ else.

Proof. $i(0) > 0$ otherwise there is no problem to solve. Observe that Ω is positively invariant. Therefore a solution starting from any initial point $\mathbf{x}(0) \in \Omega$ remains confined to Ω . This allows us to disregard state constraints from the control formulation.

Let $u(t), v(t) \in [0, 1]$ for all $t \in [0, T]$ (This relaxation allows us to establish existence of an optimal control. We show that the optimal controls are indeed ‘bang-bang’, i.e., $u^*(t), v^*(t) \in \{0, 1\}$ for all t). Writing the problem in Mayer form, it can be seen that the state space (appropriately expanded with additional variables) is bounded and positively invariant (thus, state trajectories remain

bounded for all admissible pairs); and the system is affine in controls (see (7), (4) and (5)). Existence of an optimal control is now established by Filippov-Cesari theorem.

From (4), (5), and (7), the Hamiltonian is written as follows.

$$\begin{aligned} H(\mathbf{x}, \mathbf{p}, u, v) = & -cu(\beta + \epsilon_1)ir - c'v(\alpha + \epsilon_2)i \\ & - p_1[(\beta + ue_1)ir + (\alpha + ve_2)i + \gamma i\theta + \delta i] \\ & + p_2[(\beta + ue_1)ir + (\alpha + ve_2)i] \end{aligned} \quad (8)$$

$\mathbf{p} := (p_1, p_2)$ denotes co-state variables. Then according to Pontryagin's Maximum Principle, there exist continuous and piecewise continuously differentiable co-state functions p_1 and p_2 that satisfy

$$\begin{aligned} \dot{p}_1 &= -\frac{\partial H}{\partial i} \\ &= [c\beta - (p_2 - p_1 - c)\epsilon_1]ru + [c'\alpha - (p_2 - p_1 - c')\epsilon_2]v \\ &\quad + (p_1 - p_2)(\beta r + \alpha) + p_1(\gamma(1 - 2i - r) + \delta) \end{aligned} \quad (9)$$

$$\begin{aligned} \dot{p}_2 &= -\frac{\partial H}{\partial r} \\ &= [c\beta - (p_2 - p_1 - c)\epsilon_1]ui + (p_1 - p_2)\beta i - p_1\gamma i, \end{aligned} \quad (10)$$

at all $t \in [0, T]$ where u and v are continuous and satisfy the following transversality condition

$$p_1^*(T) = 0, \quad p_2^*(T) = 1. \quad (11)$$

and also satisfy, for all $t \in [0, T]$, $u(t) \in [0, 1]$ and $v(t) \in [0, 1]$,

$$\begin{aligned} H(x^*(t), p^*(t), u^*(t), v(t)) &\geq H(x^*(t), p^*(t), u(t), v(t)) \\ H(x^*(t), p^*(t), u(t), v^*(t)) &\geq H(x^*(t), p^*(t), u(t), v(t)). \end{aligned} \quad (12)$$

From (8) and (12), we get the following form for controls.

$$u^*(t) = \begin{cases} 1 & \text{if } (p_2^*(t) - p_1^*(t) - c)\epsilon_1 > c\beta \\ 0 & \text{if } (p_2^*(t) - p_1^*(t) - c)\epsilon_1 < c\beta \end{cases} \quad (13)$$

$$v^*(t) = \begin{cases} 1 & \text{if } (p_2^*(t) - p_1^*(t) - c')\epsilon_2 > c'\alpha \\ 0 & \text{if } (p_2^*(t) - p_1^*(t) - c')\epsilon_2 < c'\alpha \end{cases} \quad (14)$$

In case of equality in the conditions specified in equations (13) and (14), $u^*(t)$ and $v^*(t)$ may take any arbitrary values in $[0, 1]$.

Let $\phi(t) := (p_2^*(t) - p_1^*(t) - c)\epsilon_1 - c\beta$ and $\psi(t) := (p_2^*(t) - p_1^*(t) - c')\epsilon_2 - c'\alpha$.

We denote by H_t^* the Hamiltonian along optimal state-control trajectory at time t . The following lemma proves that Hamiltonian will always remain positive.

Lemma 1. $H_t^* > 0 \forall t \in [0, T]$.

Proof. From (8) and (11), we have

$$\begin{aligned} H_T^* &= [(1 - c)\epsilon_1 - c\beta]i^*(T)r^*(T)u^*(T) \\ &\quad + [(1 - c')\epsilon_2 - c'\alpha]i^*(T)v^*(T) \\ &\quad + (\beta i^*(T)r^*(T) + \alpha i^*(T)) \end{aligned}$$

$r(t)$ is non-decreasing whereas $i(t) \downarrow 0$ and $i(t) > 0$ for all t since $i(0) > 0$. Therefore, $H_T^* > 0$. The conclusion follows by noting that the Hamiltonian is constant for autonomous systems. \square

The lemma below shows that the co-state variables remain positive for the whole duration.

Lemma 2. $p_1^*(t), p_2^*(t) > 0 \forall t \in [0, T]$.

Proof. Suppose $p_1^*(t) \leq 0$ for all t and let $p_1^*(t) = 0$ at $t = \tau$ (at least one τ exists since $p_1^*(T) = 0$). Then $p_2^*(\tau) > 0$ otherwise $H_\tau^* < 0$ since $u^*(\tau) = 0$. Observe from (9) that $\dot{p}_1 < 0$ if $p_1 = 0$. Strict inequality in (13) implies that at τ , $u^*(\cdot)$ is continuous. Therefore, $\dot{p}_1 < 0$ in the neighborhood of 0. Thus $p_1^*(t_1) \leq 0$ implies $p_1^*(t) < 0$ for all $t > t_1$ and $p_1^*(T) \neq 0$ which violates (11). It follows that $p_1^*(t) > 0$ for all $t \in [0, T]$. This in turn implies that $p_2^*(t) > 0$ for all t otherwise $H_t^* < 0$. \square

Lemma 3. $p_2^*(t) > p_1^*(t) \forall t \in [0, T]$.

Proof. Suppose not. Let $p_2^*(t) < p_1^*(t)$ at $t = \tau$. Then $\phi(\tau) < 0$ and, therefore, $u^*(\tau) = 0$. (8) then yields $H_\tau^* < 0$, a contradiction. \square

Let $\zeta(t) := (p_2^*(t) - p_1^*(t))i^*(t)$. The lemma that follows shows that $\zeta(t)$ is a decreasing function.

Lemma 4. $\dot{\zeta}(t) < 0 \forall t \in [0, T]$.

Proof. From (4), (9), and (10) we get

$$\begin{aligned} \dot{\zeta}(t) &= -[c(\epsilon_1 + \beta)u^*(t)r^*(t) + \phi(t)u^*(t)i^*(t) \\ &\quad + c'(\epsilon_2 + \alpha)v^*(t) \\ &\quad + p_2^*(t)(\gamma(1 - i^*(t) - r^*(t)) + \delta)]i^*(t) \end{aligned}$$

Lemma follows by noting that all terms inside the bracket are non-negative. \square

Now consider $\dot{\phi}(t)$. From (9), (10), and (8) we get

$$\dot{\phi}(t) = [\frac{H_t^*}{i^*(t)} - \phi(t)i^*(t)u^*(t) - (p_2^*(t) - p_1^*(t))\beta i^*(t)]\epsilon_1$$

$(p_2^*(t) - p_1^*(t))i^*(t)$ is monotonically decreasing (Lemma 4). From (4) $i^*(t) \downarrow 0$ exponentially ($i^*(t) < i^*(0)e^{-(\alpha+\delta)t}$). H_t^* is a positive constant (Lemma 1).

Assume that $\phi(t) = 0$ at three points in time τ_1, τ_2, τ_3 . Therefore, $\dot{\phi}(\tau) > 0$ for either $\tau = \tau_1$ or $\tau = \tau_2$. Without loss of generality, let us say $\dot{\phi}(\tau_2) > 0$. From the above equation it follows that $\dot{\phi}(\tau_3) > 0$ which is not feasible as $\phi(\tau_3^-) > 0$. It follows that $\phi(t) = 0$ at at most two points in time. Therefore, there exist $0 \leq \tau_1 \leq \tau_2 \leq T$ such that $u^*(t) = 1$ for $0 \leq t \leq \tau_1$ and $\tau_2 < t \leq T$, and 0 elsewhere.

Similarly, one can show that there exist $0 \leq \tau_3 \leq \tau_4 \leq T$ such that $v^*(t) = 0$ for $\tau_3 < t \leq \tau_4$ and $v^*(t) = 1$ otherwise. The proposition is, thus, established. ■

Proposition 1 implies that both the referral reward and direct incentives programs are to be deployed at most twice for certain durations, one in the beginning and the other at the end. It may happen that both the durations are of length 0 which means that a program is not deployed at all. On the other hand, it could also get deployed over the complete time horizon T . This gives a simple and elegant marketing strategy which is easy to implement for the seller.

The structure of the above optimal control is quite intuitive. In the case of the referral reward program, the cost is proportional to the product of number of potential buyers and customers. Hence to keep the cost low, rewards are declared in the initial stage (when the number of customers is less) to motivate product advocates and may also be paid at the end (when the number of potential buyers is less) to acquire some additional customers.

In the case of direct incentives, the cost is proportional to the number of potential buyers. If the initial take rate for the product is less, this program may get executed at initial stages to quickly acquire customers whose social influence can be exploited in the later stages; otherwise more agents may buy competitor's product and attract other potential buyers. Towards the end of the campaign, the number of potential buyers is less; hence direct incentives may be offered to attract additional customers.

4 Numerical Results

The simple structure of optimal controls given by Proposition 1 allows one to devise incentives programs quite easily by numerical optimization of τ 's. Here we obtain an independent validation of optimal controls by discretizing (7), casting it as a nonlinear constrained optimization problem and using a gradient descent approach to find an optimal solution. (For discussion on various numerical solution techniques for such problems refer to [12]). For all our experiments, the time horizon T and discretization step-size are fixed at 10 and 0.1 respectively. The NLP formulation is not convex. Therefore, we use a multi-start mechanism to determine an optimal solution. Results are also verified using the commercial package PROPT which uses pseudospectral methods for solving such problems.

In this paper, we will primarily investigate the initial condition $i(0) = 1$. This captures the case when the seller and the competitor(s) enter the market with substitutable products at around the same time (e.g., gaming technologies) or when the seller introduces an independent product into the market (e.g., a

book). Of course, similar results can be obtained for the case where seller and/or competitor already have some presence in the market ($r(0) > 0$ and/or $\theta(0) > 0$).

We fix $\epsilon_1 = \epsilon_2 = 0.05$, $\gamma = 0.1$, $\delta = 0.1$ for all numerical studies, and consider $\beta = 0.1$, $\alpha = 0.08$, $c = 0.25$ and $c' = 0.3$ as the *base scenario* (These parameter values are arbitrary and only roughly based on some available data). The optimal marketing strategy is shown in Figure 1. It is optimal for the seller to run both the incentive programs initially for some duration, stop and then run the programs again towards the end. Note that the optimal strategy is open loop; hence estimates of i and r are not required for implementation.

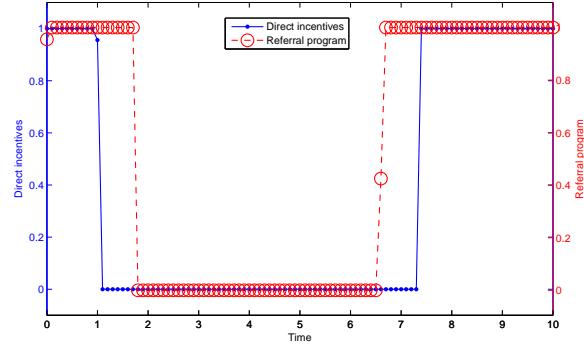


Fig. 1. Optimal marketing strategy for the base scenario

It is, thus, possible for the seller to determine the timing of her incentives programs numerically. Experimentation with different values of pay-outs c and c' (which essentially fix ϵ_1 and ϵ_2) can be used to understand trade-offs and optimize these pay-outs.

In the following we undertake an investigation of two important questions pertaining to the interplay between the two incentives programs and the impact of network structure on them. The former question is important because *influence-and-exploit* strategy has received much attention in the literature. As we show below, *exploit-and-influence* strategy can also come into play for some parameter settings. The second question is linked to the so-called *influentials hypothesis* which informally says that high degree agents (*hubs*) play significant role in product diffusion, and, therefore, are natural targets for incentives (direct or referral rewards). We show that in some cases the seller is better off incentivizing low degree agents (more than high-degree ones). Thus, our results highlight the need for a careful consideration of the network structure while making incentive decisions.

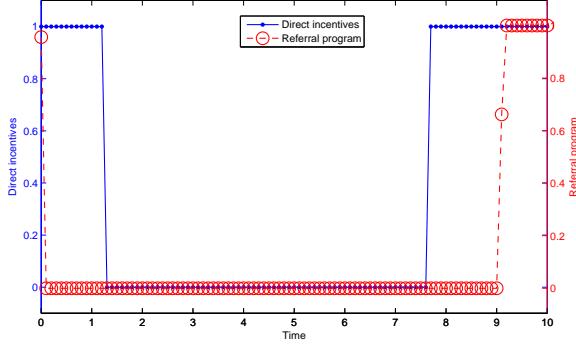


Fig. 2. *Influence-and-exploit* strategy is optimal when β is increased to 0.13

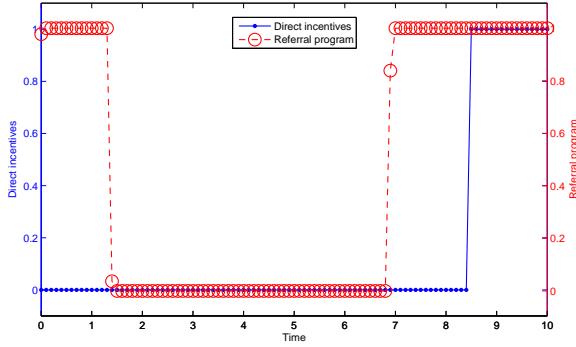


Fig. 3. *Exploit-and-influence* strategy is optimal when α is increased to 0.09

4.1 Interplay between Referral and Direct Incentive Programs

When social influence is strong in the population, i.e., β is higher, the seller needs to employ only direct incentives initially. For example, if β is set to 0.13 in the base scenario then it is optimal to offer referral rewards only at the end and that too for a short period as shown in Figure 2. This can be seen as a manifestation of the *influence-and-exploit* strategy. On the other hand, if the seller has established a good reputation in the market, translating into a higher value of α , an initial influence step through direct incentives may not even be required. See from Figure 3 that when $\alpha = 0.09$ in the base scenario, it is optimal for the seller to offer direct incentives only at the end. In this case, for most portion of the time horizon, the seller must exploit connections of existing customers and only at the end must she impart direct influence on potential buyers. We call it the *exploit-and-influence* strategy for the seller.

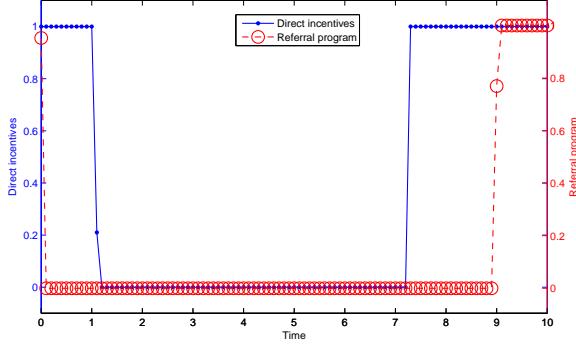


Fig. 4. *Influence-and-exploit* strategy is optimal with pay-outs: $c = 0.3$, $c' = 0.3$

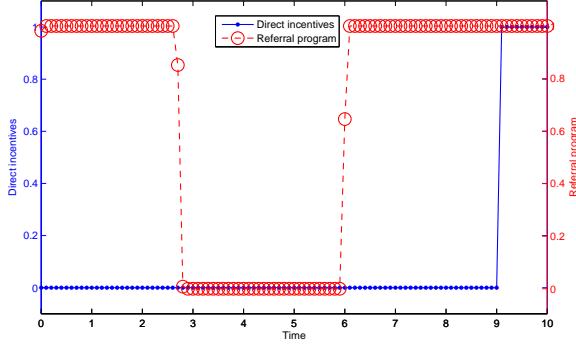


Fig. 5. *Exploit-and-influence* strategy is optimal with pay-outs: $c = 0.25$, $c' = 0.35$

We also observe from Figures 4 and 5 that *influence-and-exploit* and *exploit-and-influence* are optimal strategies for the seller if she incurs high per conversion pay-outs for referral and incentive programs respectively.

4.2 Impact of Network Structure on Incentive Programs

Real-world networks show strong degree correlation amongst connected nodes. Some networks show *assortative* mixing of nodes by degrees where high-degree nodes have most of their connections to other high-degree nodes. Others show *disassortative* mixing where high-degree nodes have most of their connections to low-degree nodes ([13]). In this section, we examine the impact of network structure on incentives programs.

We consider an undirected correlated network with nodes belonging to either of two classes A and B with probability $P(A)$ and $P(B)$ respectively. Class A

nodes are of high degree, say k_A and class B nodes are of low degree, say k_B . $P(A|B)$ is the probability that a given link from class B node points to a class A node. $P(B|A)$ can be computed from the following balance equation:

$$k_A P(B|A)P(A) = k_B P(A|B)P(B) \quad (15)$$

We consider two types of network structures. One structure represents assortative mixing whereas the other one represents disassortative mixing. To keep things simple and derive key insights, we assume that the seller is optimizing implementation of only one incentives program, namely, referral rewards program. The seller can offer referral rewards to class A and/or class B nodes to increase their social influence rate (β) by ϵ . As earlier, she incurs a per conversion cost of c after normalizing with respect to product price.

We fix $\alpha = 0.1$, $\delta = 0.1$, $\beta = 0.1$, $\gamma = 0.15$, $\epsilon = 0.08$, $k_A = 10$, $k_B = 2$, $P(A) = 0.1$, $P(B) = 0.9$ for our numerical studies. For disassortative network, we set $P(B|A) = 0.9$ whereas for assortative network, we set $P(B|A) = 0.1$.

The optimal timing of referral reward program for disassortative network is shown in Figure 6. The seller's optimal strategy is to offer referral rewards to class B nodes for the complete duration whereas rewards to class A nodes are offered initially for a short duration and then again towards the end for a short duration. In this case, class B nodes have almost half of their connections going to class A nodes. Also, major fraction of the population is from class B . So, referral rewards are offered to class B nodes for entire duration as it increases influence not only on class B nodes but also on class A nodes. Class A nodes are not rewarded for the entire duration in order to control the cost.

In the case of assortative network, the optimal strategy changes completely (see Figure 7). The seller offers referral rewards to class A nodes for the complete duration whereas rewards to class B nodes are offered initially for some duration and then again towards the end. In this case, nodes from both the classes are well connected amongst themselves with very few connections going across the classes. So, the optimal reward strategies for both the classes are essentially independent. For this particular scenario, it turns out that it is optimal to offer rewards to class A nodes for the entire duration as the cost incurred is not much. Whereas in the case of class B nodes, referral rewards are discontinued for some duration in the middle as the cost overshoots the potential revenue.

The results show that networks with different structures can result in different optimal strategies for the seller. In some scenarios, the seller may be better off incentivizing low degree nodes as against the popular approach of targeting the influentials (high degree nodes), thus, providing some support to the finding in [14]. In some scenarios, the seller may be better off targeting influentials thus supporting the results in [9]. Thus, our results highlight the need for a careful consideration of the network structure while making incentive decisions.

5 Conclusion

In this paper we have addressed the problem of optimal timing of two incentive programs, namely, direct incentives and referral rewards, for product diffusion

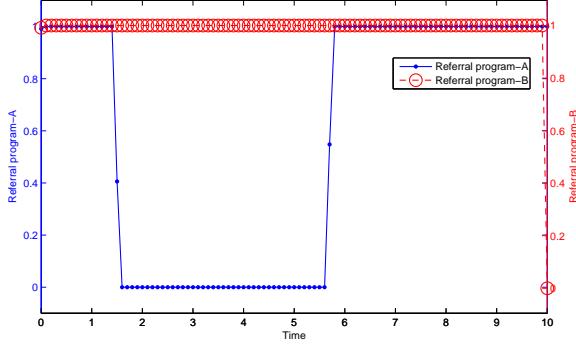


Fig. 6. Optimal marketing strategy on a network with disassortative mixing

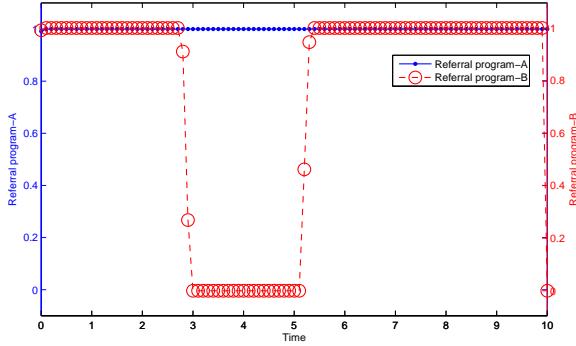


Fig. 7. Optimal marketing strategy on a network with assortative mixing

through social networks. Taking a deviation from the existing approaches, we formulate the problem as a continuous-time deterministic optimal control problem. The optimal strategy for the seller is to deploy these programs in at most two distinct time periods. The simplicity of this structure and non-adaptive nature makes them ideal for implementation in practice. We further show that if the seller has good reputation in the market, *exploit-and-influence* strategy can be optimal whereas if social influence is strong in the population, *influence-and-exploit* strategy can be optimal for the seller. In the case of correlated networks, our numerical studies show that the seller need not necessarily offer more frequent referral reward programs to high degree nodes to maximize her profit.

There are two immediate directions for future work: (i) extend heterogeneity of agents to include their external and social influence probabilities and (ii) devise procedures to estimate model parameters.

References

- [1] D. Arthur, R. Motwani, A. Sharma, and Y. Xu. Pricing strategies for viral marketing on Social Networks. In *Proceedings of Workshop on Internet and Network Economics*, pages 101–112, 2009.
- [2] E. Auriol and M. Benaim. Standardization in decentralized economies. *American Economic Review*, 90(3):550–570, 2000.
- [3] M. Benaim and J. Y. Le Boudec. A Class of Mean Field Interaction Models for Computer and Communication Systems. *Performance Evaluation*, 65: 823–838, 2008.
- [4] S. Bharathi, D. Kempe, and M. Salek. Competitive influence maximization in social networks. In *Proceedings of Workshop on Internet and Network Economics*, pages 306–311, 2007.
- [5] C. Budak, D. Agrawal, and A. El Abbadi. Limiting the Spread of Misinformation in Social Networks. In *Proceedings of the 20th International Conference on World Wide Web*, pages 665–674, 2011.
- [6] W. Chen, Y. Yaun, and L. Zhang. Scalable influence maximization in social networks under the linear threshold model. In *Proceedings of 10th International Conference on Data Mining*, 2010.
- [7] P. Domingos and M. Richardson. Mining the Network Value of Customers. In *7th International Conference on Knowledge Discovery and Data Mining*, pages 57–66, 2001.
- [8] D. Godes, D. Mayzlin, Y. Chen, S. Das, C. Dellarocas, B. Pfeiffer, B. Libai, S. Sen, M. Shi, and P. Verlegh. The Firm’s Management of Social Interactions. *Marketing Letters*, 16(3/4):415–428, 2005.
- [9] J. Goldenberg, S. Han, D. R. Lehmann, and J. W. Hong. The Role of Hubs in the Adoption Processes. *Journal of Marketing*, 73:1–13, March 2009.
- [10] J. Hartline, V. S. Mirrokni, and M. Sundararajan. Optimal Marketing Strategies over Social Networks. In *Proceedings of 17th International Conference on World Wide Web*, pages 189–198, 2008.
- [11] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the Spread of Influence through a Social Network. In *Proceedings of Knowledge Discovery and Data Mining*, pages 137–146, 2003.
- [12] D. E. Kirk. *Optimal Control Theory: An Introduction*. Dover, 1970.
- [13] M. E. J. Newman. Assortative Mixing in Networks. *Physics Review Letters*, 89(20), November 2002.
- [14] D. Watts and P. S. Dodds. Influentials, Networks, and Public Opinion Formation. *Journal of Consumer Research*, 34, December 2007.